



US009160805B2

(12) **United States Patent**  
**Yoo et al.**

(10) **Patent No.:** **US 9,160,805 B2**  
(45) **Date of Patent:** **\*Oct. 13, 2015**

(54) **PROXY-BASED CACHE CONTENT DISTRIBUTION AND AFFINITY**

(58) **Field of Classification Search**  
USPC ..... 709/213, 217, 218, 219  
See application file for complete search history.

(71) Applicant: **MICROSOFT CORPORATION**,  
Redmond, WA (US)

(56) **References Cited**

(72) Inventors: **Won Suk Yoo**, Redmond, WA (US);  
**Nick Holt**, Seattle, WA (US); **Daniel Vasquez Lopez**, Duvall, WA (US);  
**Aniello Scotto Di Marco**, Redmond, WA (US)

U.S. PATENT DOCUMENTS

5,305,389 A 4/1994 Palmer  
5,924,116 A 7/1999 Aggarwal  
6,330,606 B1 12/2001 Logue et al.  
6,427,187 B2 7/2002 Malcolm  
6,438,652 B1 8/2002 Jordan et al.

(Continued)

(73) Assignee: **Microsoft Technology Licensing, LLC**,  
Redmond, WA (US)

OTHER PUBLICATIONS

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

Das, et al., "Predicting Web Cache Behavior using Stochastic State-Space Models", In Proceedings of the Second International Workshop on Scalable Data Management Applications and Systems, Jul. 2008, 8 pages. Available at <<http://csl.cse.psu.edu/publications/cachepred-sdmas08.pdf>>.

(Continued)

(21) Appl. No.: **14/096,697**

(22) Filed: **Dec. 4, 2013**

*Primary Examiner* — Umar Cheema

*Assistant Examiner* — Marshall McLeod

(65) **Prior Publication Data**

US 2014/0095649 A1 Apr. 3, 2014

(74) *Attorney, Agent, or Firm* — Ben Tabor; Doug Barker; Micky Minhas

**Related U.S. Application Data**

(63) Continuation of application No. 13/022,527, filed on Feb. 7, 2011, now Pat. No. 8,612,550.

(57) **ABSTRACT**

(51) **Int. Cl.**

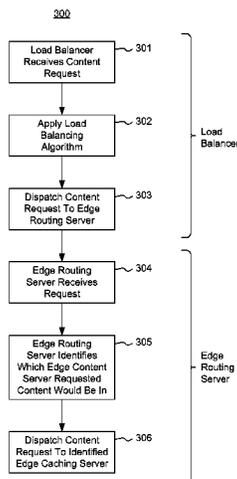
**G06F 15/167** (2006.01)  
**H04L 29/08** (2006.01)  
**H04L 12/721** (2013.01)  
**H04L 12/775** (2013.01)

A distributed caching hierarchy that includes multiple edge routing servers, at least some of which receiving content requests from client computing systems via a load balancer. When receiving a content request, an edge routing server identifies which of the edge caching servers the requested content would be in if the requested content were to be cached within the edge caching servers, and distributes the content request to the identified edge caching server in a deterministic and predictable manner to increase the likelihood of increasing a cache-hit ratio.

(52) **U.S. Cl.**

CPC ..... **H04L 67/2842** (2013.01); **H04L 45/44** (2013.01); **H04L 45/58** (2013.01); **H04L 67/288** (2013.01); **H04L 67/2885** (2013.01); **H04L 67/327** (2013.01)

**8 Claims, 4 Drawing Sheets**



(56)

**References Cited**

U.S. PATENT DOCUMENTS

6,490,615	B1	12/2002	Dias et al.	
6,615,317	B2	9/2003	Roseborough et al.	
6,732,237	B1	5/2004	Jacobs	
7,024,466	B2	4/2006	Outten et al.	
7,133,905	B2	11/2006	Dilley et al.	
7,565,423	B1	7/2009	Fredricksen	
7,697,557	B2	4/2010	Segel	
7,730,262	B2	6/2010	Lowery et al.	
7,761,664	B2	7/2010	Gill	
7,904,562	B2	3/2011	Takase et al.	
8,291,007	B2	10/2012	Orr et al.	
2004/0167981	A1	8/2004	Douglas	
2004/0172466	A1	9/2004	Douglas	
2004/0254943	A1	12/2004	Malcolm	
2008/0222281	A1*	9/2008	Dilley et al. ....	709/223
2009/0248871	A1	10/2009	Takase et al.	
2010/0080226	A1	4/2010	Khalid et al.	
2010/0332595	A1*	12/2010	Fullagar et al. ....	709/203
2012/0185557	A1	7/2012	Di Marco et al.	
2012/0203866	A1	8/2012	Yoo et al.	

OTHER PUBLICATIONS

Gill, Binny S., "On Multi-Level Excluding Caching: Offline Optimality and Why promotions are better than demotions.", In 6th USENIX

Conference on File and Storage Technologies, Feb. 27, 2008, 17 pages. Available at <<[http://webcache.googleusercontent.com/search?q=cache:nwgGP4zklyAJ:static.usenix.org/event/fast08/tech/full\\_papers/gill/gill.pdf+&cd=1&hl=en&ct=clnk&gl=in](http://webcache.googleusercontent.com/search?q=cache:nwgGP4zklyAJ:static.usenix.org/event/fast08/tech/full_papers/gill/gill.pdf+&cd=1&hl=en&ct=clnk&gl=in)>>.

Gulwani et al., WebCaL—A Domain Specific Language for Web Caching, May 2000, [http://research.microsoft.com/en-us/um/people/sumitg/pubs/webcal\\_wcw00.pdf](http://research.microsoft.com/en-us/um/people/sumitg/pubs/webcal_wcw00.pdf).

Lindemann, "Evaluating Cooperative Web Caching Protocols for Emerging Network Technologies," University of Dortmund, Department of Computer Science, Aug. 2001.

Andrey Naumenko, "Some Recommendations on Building Proxy caching Service", Pub date: Nov. 1998 (6 pages).

Duane Wessels et al., "ICP and the Squid Web Cache", Pub date : Aug. 13, 1997 (25 pages).

Nikolaos Laoutaris, et al., "The Cache Inference Problem and its Application to Content and Request Routing", Pub. Date: May 29, 2007 (10 pages).

U.S. Appl. No. 13/006,694, Aug. 29, 2012, Office Action.

U.S. Appl. No. 13/006,694, Nov. 20, 2012, Notice of Allowance.

U.S. Appl. No. 13/022,527, Mar. 22, 2013, Office Action.

U.S. Appl. No. 13/022,527, Aug. 20, 2013, Notice of Allowance.

\* cited by examiner

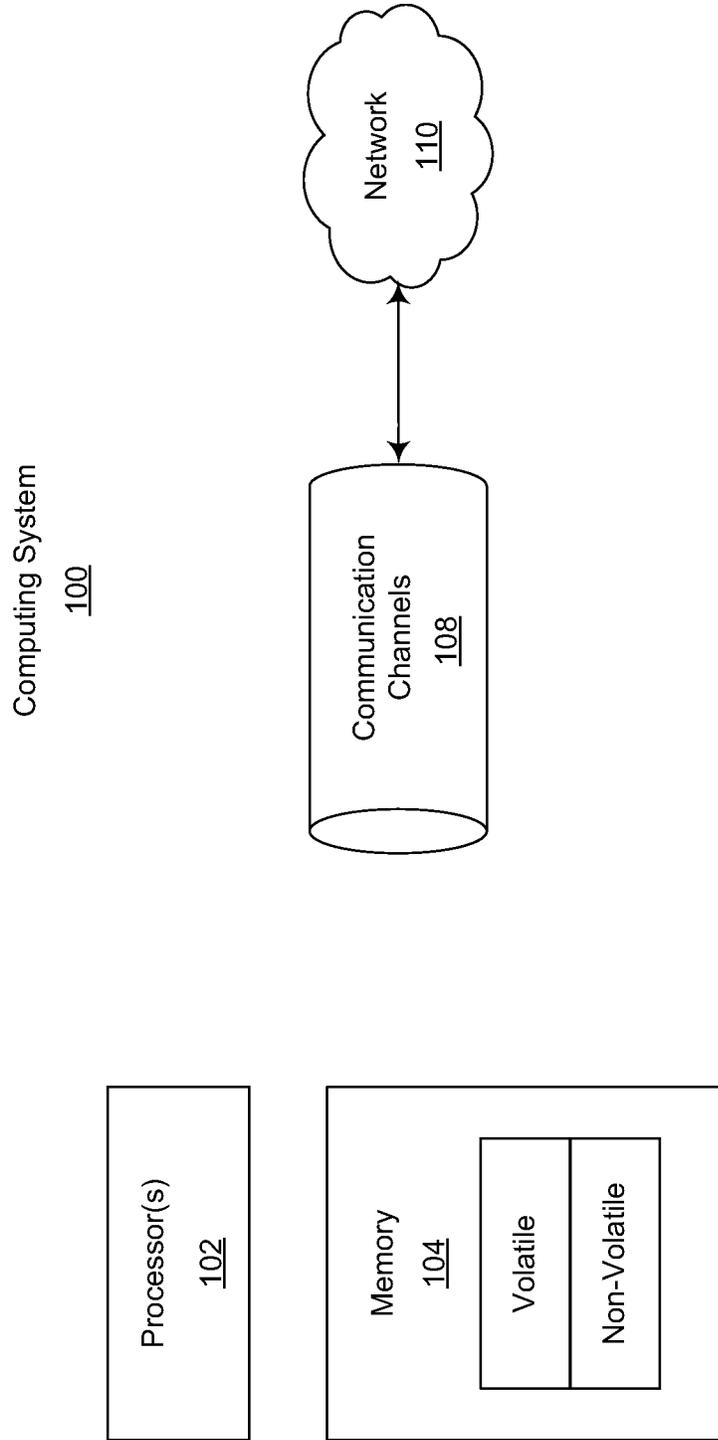


Figure 1

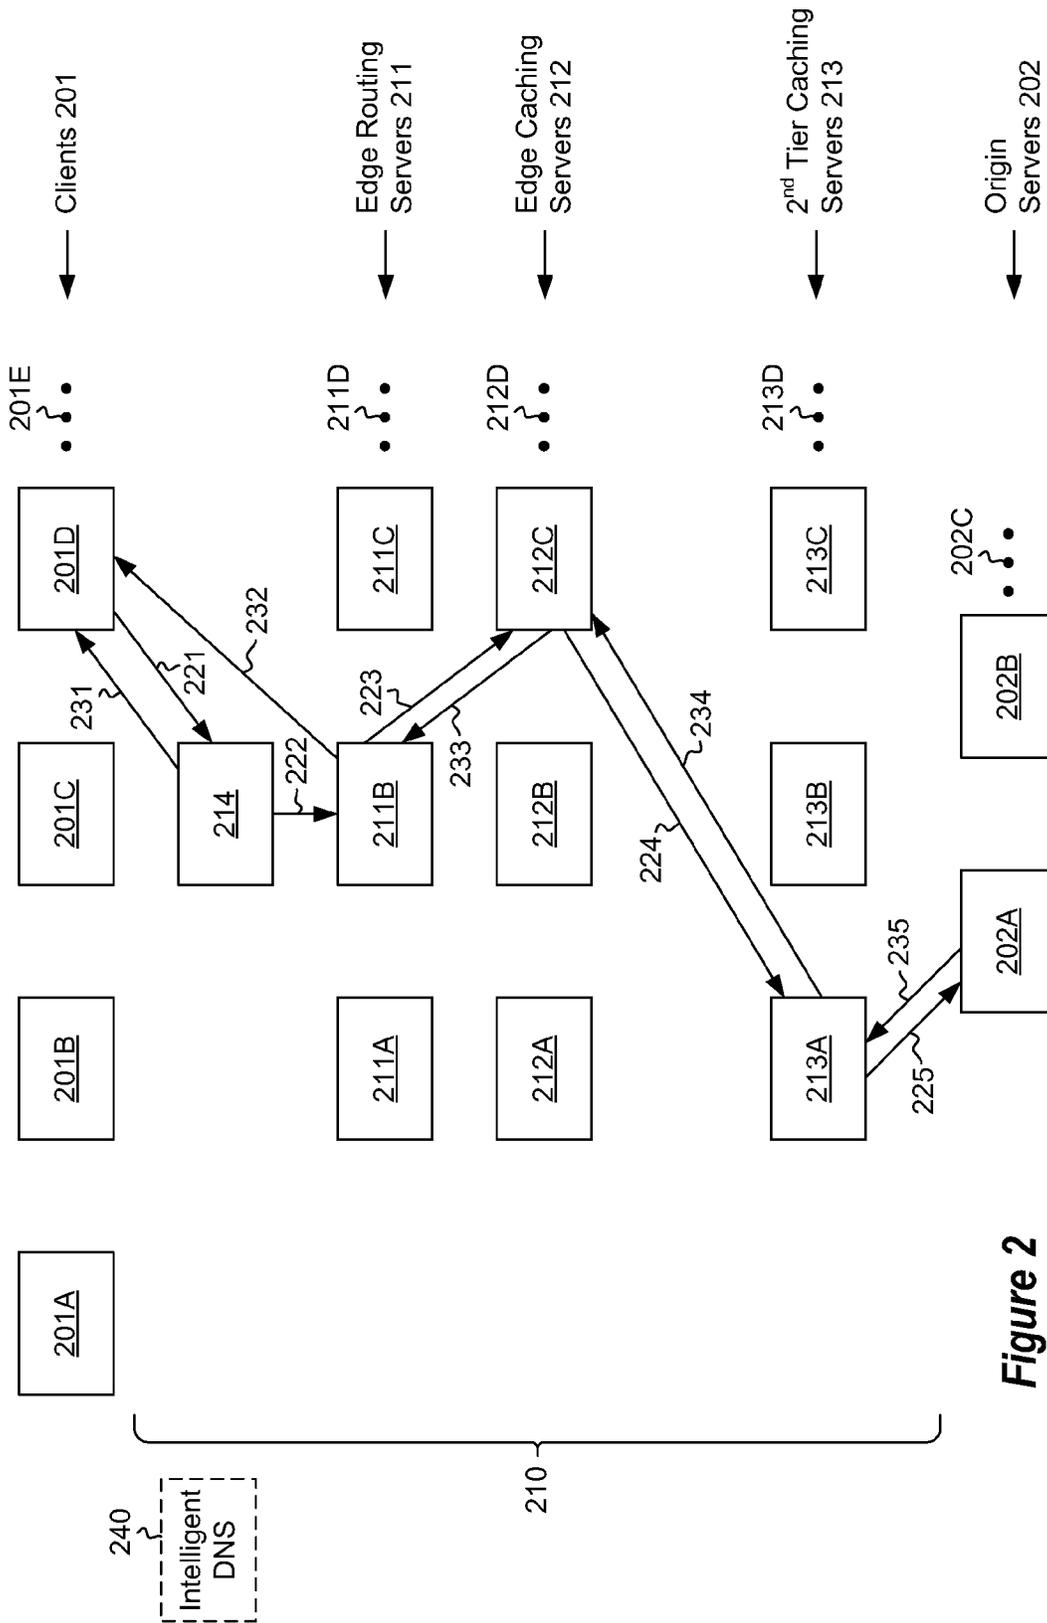
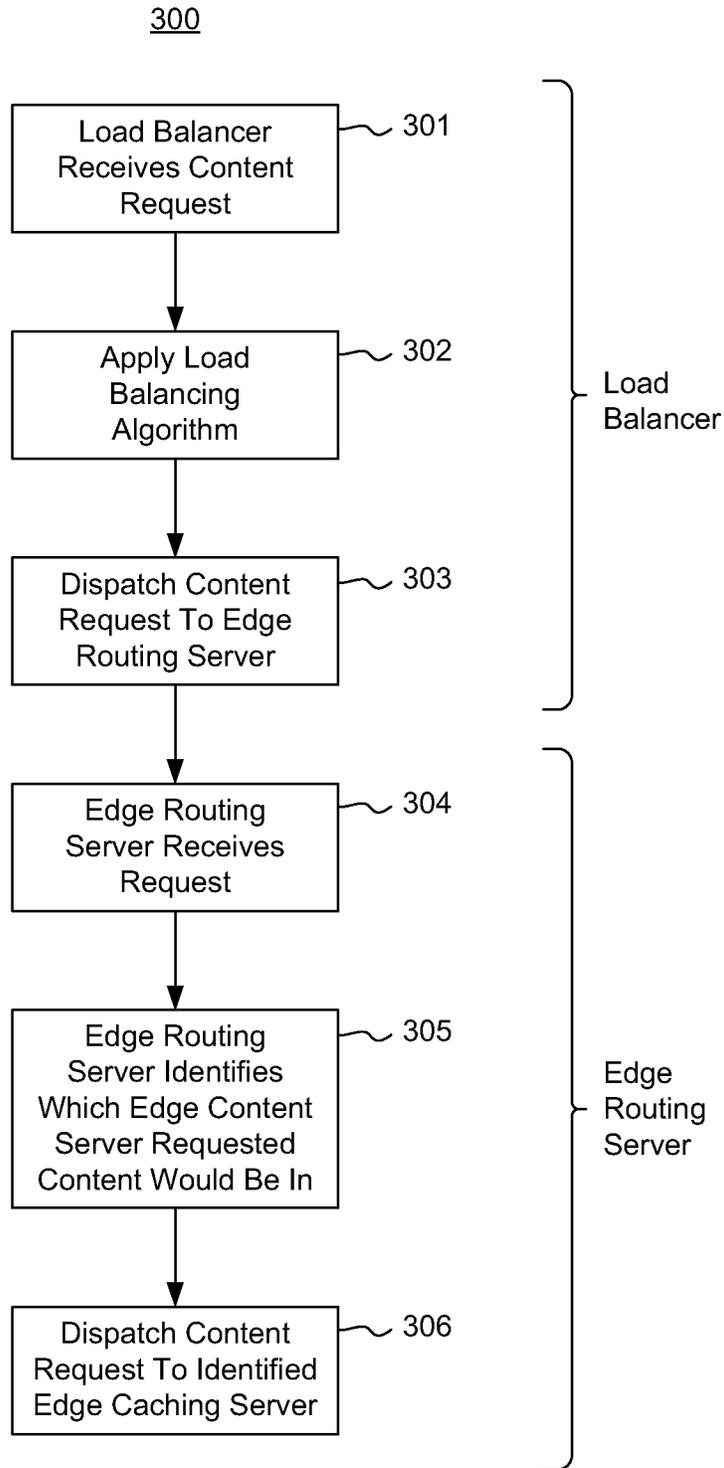


Figure 2



**Figure 3**

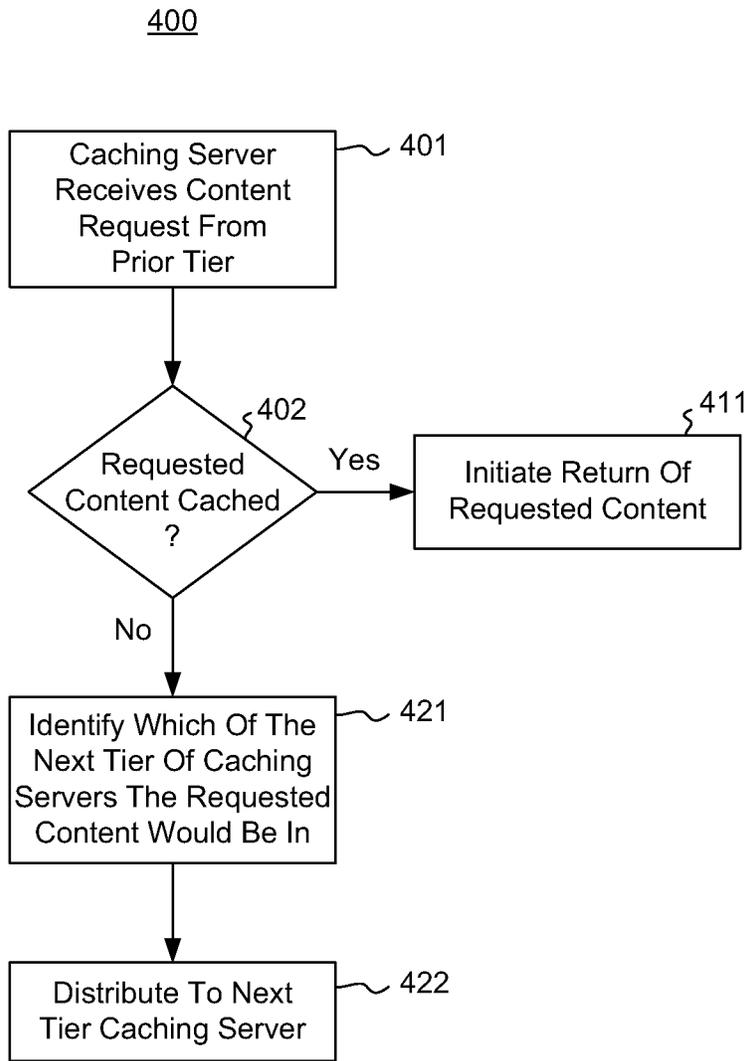


Figure 4

**PROXY-BASED CACHE CONTENT  
DISTRIBUTION AND AFFINITY****CROSS-REFERENCE TO RELATED  
APPLICATIONS**

This application is a continuation of U.S. patent application Ser. No. 13/022,527 filed on Feb. 7, 2011, entitled “PROXY-BASED CACHE CONTENT DISTRIBUTION AND AFFINITY,” which issued as U.S. Pat. No. 8,612,550 on Dec. 17, 2013, and which application is expressly incorporated herein by reference in its entirety.

**BACKGROUND**

Content distribution networks are used to deliver all types of content from origin servers to client computing systems. To efficiently deliver such content, distributed cache hierarchies intervene between origin servers and client computing systems. When a client issues a request for content, if the content is available in the hierarchy, the content may often be provided by the hierarchy rather than having to obtain the content all the way from the origin server, provided that the content is found in the hierarchy.

In many cases, the cache hierarchy is partitioned in such a way that content items are cached mutually exclusively with one another, while providing a way to route the requests accordingly and expectedly to find the cached content. Therefore, when a request is routed to a cache server, the cache content is expected to be cached by that cache server and it would likely result in a cache hit, meaning the requested content is cached by that cache server and thus the cache server may satisfy the content request directly. One way to do this is to perform such cache distribution and predictive routing at the edge caching tier (in a cache hierarchy, this would be the first tier of cache servers geographically closest to the client).

A conventional way for a client to determine a location of cached content in an edge caching tier is for the client to first provide the Uniform Resource Locator (URL) of the desired content to an intelligent Domain Name Server (DNS) server. The intelligent DNS server knows the geographic proximity of the cached content, and provides an Internet Protocol (IP) address of a cache server that would perform well for the client. The client then requests the content using the IP address provided by the intelligent DNS server.

**BRIEF SUMMARY**

At least one embodiment described herein relates to a distributed caching hierarchy that includes multiple edge routing servers, at least some of which receive content requests from client computing systems via a load balancer. When receiving a content request, an edge routing server identifies which of the edge caching servers is responsible for caching the requested content in a predictable manner. Once the appropriate caching server has been identified, the edge routing server routes the content request to the appropriate caching server to maximize the likelihood of a cache hit. The edge routing server and the edge caching server can be in one physical server in a cache hierarchy, although that is not required.

In accordance with at least one embodiment described herein, when the edge caching server receives the content request, the edge caching server determines whether the edge caching server has the content cached. If so, the edge caching server initiates return of the requested content to the client. If

not, the edge caching server identifies which of the next tier caching servers is responsible for the requested content in a predictable manner and routes the now cache-missed content request to the appropriate next tier caching server. This continues potentially all the way up the distributed caching hierarchy until, if necessary, the origin server is reached. When the next tier caching server receives the requested content, it caches the requested content, and then initiates return of the requested content to the client.

This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

**BRIEF DESCRIPTION OF THE DRAWINGS**

In order to describe the manner in which the above-recited and other advantages and features can be obtained, a more particular description of various embodiments will be rendered by reference to the appended drawings. Understanding that these drawings depict only sample embodiments and are not therefore to be considered to be limiting of the scope of the invention, the embodiments will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

FIG. 1 illustrates an example computing system that may be used to employ embodiments described herein;

FIG. 2 illustrates a network that includes a distributed cache hierarchy;

FIG. 3 illustrates a flowchart of a method for managing a distributed cache hierarchy in a network from the viewpoint of a load balancer and a tier of edge routing servers; and

FIG. 4 illustrates a flowchart of a method for managing a distributed cache hierarchy in a network from the viewpoint of a tier of edge caching servers.

**DETAILED DESCRIPTION**

In accordance with embodiments described herein, a distributed caching hierarchy includes multiple edge routing servers, at least some of which receive content requests from client computing systems via a load balancer. When receiving a content request, an edge routing server identifies which of the edge caching servers is responsible for the requested content in a predictable manner and routes the now cache-missed content request to the identified next tier caching server. First, some introductory discussion regarding computing systems will be described with respect to FIG. 1. Then, the embodiments of the distributed caching network and its example operation will be described with respect to subsequent figures.

First, introductory discussion regarding computing systems is described with respect to FIG. 1. Computing systems are now increasingly taking a wide variety of forms. Computing systems may, for example, be handheld devices, appliances, laptop computers, desktop computers, mainframes, distributed computing systems, or even devices that have not conventionally been considered a computing system. In this description and in the claims, the term “computing system” is defined broadly as including any device or system (or combination thereof) that includes at least one physical and tangible processor, and a physical and tangible memory capable of having thereon computer-executable instructions that may be executed by the processor. The memory may take any form and may depend on the nature and form of the computing system. A computing system may be distributed over a network environment and may include multiple constituent com-

puting systems. As illustrated in FIG. 1, in its most basic configuration, a computing system **100** typically includes at least one processing unit **102** and memory **104**. The memory **104** may be physical system memory, which may be volatile, non-volatile, or some combination of the two. The term “memory” may also be used herein to refer to non-volatile mass storage such as physical storage media. If the computing system is distributed, the processing, memory and/or storage capability may be distributed as well. As used herein, the term “module” or “component” can refer to software objects or routines that execute on the computing system. The different components, modules, engines, and services described herein may be implemented as objects or processes that execute on the computing system (e.g., as separate threads).

In the description that follows, embodiments are described with reference to acts that are performed by one or more computing systems. If such acts are implemented in software, one or more processors of the associated computing system that performs the act direct the operation of the computing system in response to having executed computer-executable instructions. An example of such an operation involves the manipulation of data. The computer-executable instructions (and the manipulated data) may be stored in the memory **104** of the computing system **100**. Computing system **100** may also contain communication channels **108** that allow the computing system **100** to communicate with other message processors over, for example, network **110**.

Embodiments of the present invention may comprise or utilize a special purpose or general-purpose computer including computer hardware, such as, for example, one or more processors and system memory, as discussed in greater detail below. Embodiments within the scope of the present invention also include physical and other computer-readable media for carrying or storing computer-executable instructions and/or data structures. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer system. Computer-readable media that store computer-executable instructions are physical storage media. Computer-readable media that carry computer-executable instructions are transmission media. Thus, by way of example, and not limitation, embodiments of the invention can comprise at least two distinctly different kinds of computer-readable media: computer storage media and transmission media.

Computer storage media includes RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer.

A “network” is defined as one or more data links that enable the transport of electronic data between computer systems and/or modules and/or other electronic devices. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a transmission medium. Transmission media can include a network and/or data links which can be used to carry or desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer. Combinations of the above should also be included within the scope of computer-readable media.

Further, upon reaching various computer system components, program code means in the form of computer-execut-

able instructions or data structures can be transferred automatically from transmission media to computer storage media (or vice versa). For example, computer-executable instructions or data structures received over a network or data link can be buffered in RAM within a network interface module (e.g., a “NIC”), and then eventually transferred to computer system RAM and/or to less volatile computer storage media at a computer system. Thus, it should be understood that computer storage media can be included in computer system components that also (or even primarily) utilize transmission media.

Computer-executable instructions comprise, for example, instructions and data which, when executed at a processor, cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. The computer executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, or even source code. Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the described features or acts described above. Rather, the described features and acts are disclosed as example forms of implementing the claims.

Those skilled in the art will appreciate that the invention may be practiced in network computing environments with many types of computer system configurations, including, personal computers, desktop computers, laptop computers, message processors, hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, mobile telephones, PDAs, pagers, routers, switches, and the like. The invention may also be practiced in distributed system environments where local and remote computer systems, which are linked (either by hardwired data links, wireless data links, or by a combination of hardwired and wireless data links) through a network, both perform tasks. In a distributed system environment, program modules may be located in both local and remote memory storage devices.

FIG. 2 illustrates a network that includes a distributed cache hierarchy. The network **200** includes a number of client computing systems **201A**, **201B**, **201C** and **201D** (collectively referred to as clients **201**) and a plurality of origin server computing systems **202A** and **202B** (collectively referred to as origin servers **202**) with a cache hierarchy **210** disposed therebetween.

The “origin” server is the server that originally sources the content. Although content is available on the origin servers **202**, that content may often be available within the cache hierarchy **210** when requested by one of the clients **201**. The ellipsis **201E** represents that there may be any number of clients within the network, and that the number of clients need not be static. The ellipsis **202C** represents that there may be any number of origin servers within the network, and again that the number of origin servers need not be static. Each of the clients **201** and origin servers **202** may be structured as described above for the computing system **100**.

The example cache hierarchy **210** includes a tier of edge routing servers **211A**, **211B** and **211C** (collectively referred to as edge routing servers **211**), a tier of edge caching servers **212A**, **212B** and **212C** (collectively referred to as edge caching servers **212**) and a second tier of caching servers **213A**, **213B**, and **213C** (collectively referred to as second tier caching servers **213**). The edge routing servers **211** may perhaps not actually cache content, but instead they identify the edge cache server for the requested content in a predictable manner

to increase the likelihood of a cache-hit. In addition, the edge routing server and the edge caching server are logical roles. These roles could be served by the same physical server in the hierarchy. The edge cache servers and the second tier cache servers actually cache content.

A load balancer **214** receives content requests from the clients **201** and distributes the content request to one of the edge routing servers **211** in accordance with a load balancing algorithm. The load balancer **214** may also monitor a status of edge routing servers **211** and avoid distributed content requests to edge routing servers that are not healthy enough to sufficiently operate on such requests. The requested content may be video, images, files, sound, music, collaborative authoring content, or any other content.

While the load balancer is a good solution to provide both high availability and scalability, it does not have any awareness of where in the cache hierarchy the requested content resides. This is inefficient when it comes to routing requests for cached contents because its load balancing algorithm does not help to maximize the likelihood of getting a cache hit. The principles described herein reduce the impact of this inefficiency since the edge routing servers more predictably route the content requests to the edge caching servers. Because the edge routing servers can determine which edge caching server is responsible for caching the requested content, by routing the request to the responsible server, it can increase the cache-hit ratio.

The example cache hierarchy **210** is kept simple for clarity in explaining the principles described herein. An actual cache hierarchy may have any number of tiers of cache servers intermediating between the origin servers **202** and the clients **201**. The cache hierarchy **210** may, in fact, be distributed across the Internet. Furthermore, the ellipses **211D**, **212D** and **213D** represent that there may be any number of edge routing servers, edge caching servers, and second tier caching servers. The number of edge routing servers need not be the same as the number of caching servers in any tier, and vice versa. In one embodiment, at least one, but potentially all, of the edge routing servers may be on the same machine as a respective edge caching server. For instance, edge routing server **211A** may be on the same machine as edge caching server **212A**, edge routing server **211B** may be on the same machine as edge caching server **212B**, and/or edge routing server **211C** may be on the same machine as edge caching server **213C**, although not required. Each of the servers in the distributed cache hierarchy **210** may be structured as described above for the computing system **100** of FIG. **1**.

FIG. **3** illustrates a flowchart of a method **300** for managing a distributed cache hierarchy in a network. While the method **300** may be performed in any network having a tier of edge routing servers and at least one tier of edge caching servers, the method **300** will now be described with respect to the network **200** of FIG. **2**. Specifically, the arrows in FIG. **2** represent the communications associated with a specific content request.

The load balancer receives the content request (act **301**), and applies a load balancing algorithm to determine which edge routing server should handle the content request (act **302**). The load balancer then distributes the content request to the corresponding edge routing server (act **303**). The precise load balancing algorithm is not critical to the principles described herein and may be random, round robin, or any other load balancing algorithm because these algorithms do not have awareness of which server is responsible for caching particular content. Referring to FIG. **2**, the load balancer **214** receives the content requested as represented by arrow **221**. The load balancer **214** determines the edge routing server to

which the content request is to be routed. In the illustrated case, the load balancer **214** routes the content request to the edge routing server **211B** as represented by arrow **222**.

The edge routing server receives the content request from the load balancer (act **304**) (also as represented by the arrow **222**), the edge routing server identifies which of the edge caching servers the requested content would be in if the requested content were to be cached within the edge caching servers (act **305**). For example, the edge routing server **211B** may apply the Cache Array Routing Protocol (CARP) using the Uniform Resource Locator (URL) to identify the corresponding edge caching server. The edge routing server does not itself cache the requested content. Accordingly, the edge routing server dispatches the content request to the identified edge caching server (act **306**). Referring to FIG. **2**, for example, the edge routing server **211B** routes the content request to the corresponding edge caching server **212C** as represented by arrow **223**. Any of the other edge routing servers **211** of FIG. **2** may behave the same way with respect to routing to edge caching servers **212** that are appropriate given the URL of the content request.

In one embodiment, in addition to routing the requests in a predictable manner, the edge routing servers also route the requests in a way that the cached contents on the edge cache servers are mutually exclusive. For instance, if edge caching server **212A** is responsible for caching a particular content, then the same content is not cached by another other edge caching server **212**. This behavior allows the cache hierarchy operator to distribute the contents in a way that there are no duplicated cached contents in the cache hierarchy, and therefore to more efficiently use resources, which would have otherwise been wasted by caching duplicate contents.

FIG. **4** illustrates a flowchart of a method **400** for a caching server at any tier in the distributed cache hierarchy to handle requests from the prior tier in the distributed cache hierarchy. For instance, if the caching server was an edge caching server amongst the edge caching servers **212**, the edge caching server would perform the method **400** in response to receiving a content request from an edge routing server of the edge routing servers **211**. If the caching server was a second tier caching server of the second tier caching servers **213**, the second tier caching server would perform the method **400** in response to receiving the content request from an edge caching server of the edge content servers **212**, and so forth for subsequent tiers.

Applying method **400** first to the edge caching server tier, the edge caching server receives the content request from prior tier in the distributed cache hierarchy (act **401**), which would be the edge routing server in the case where the method **400** is applied to the edge caching server tier. For instance, in the example of FIG. **2**, the edge caching server **212C** receives the content request from the edge routing server **211B** as also represented by arrow **223**.

The edge caching server determines whether the requested content is cached on the edge caching server (decision block **402**). If the requested content is cached in the edge caching server (Yes in decision block **402**), the edge caching server initiates return of the requested content (act **411**) to a client computing system associated with the content request. Referring to FIG. **2**, if the edge caching server **212C** had the requested content, the requested content would be returned along the reverse path along arrow **233** and arrow **232** in sequence to the client **201D**.

If the requested content is not cached in the edge caching server (No in decision block **402**), the edge caching server identifies which of the next tier of caching servers the requested content would be in (act **421**) if the requested

content were to be cached within the next tier caching servers. The content request is then distributed to the identified next tier caching server (act 422). For instance, if the requested content were not in the edge caching server 212C, the content request would be dispatched to the identified next tier caching server. In the illustrated case, the CARP algorithm may be applied by the edge caching server to identify that the requested content would be cached by the second tier caching server 213A if the requested content were cached by one of the second tier caching servers 213. The method 400 would then be performed for the next tier caching server.

Applying method 400 next to the second tier caching servers, the second tier caching server receives the content request from prior tier in the distributed cache hierarchy (act 401), which would be the edge caching server in the case where the method 400 is applied to the second tier caching server tier. For instance, in the example of FIG. 2, the second tier caching server 213A receives the content request from the edge caching server 212C as also represented by arrow 224.

The second tier caching server determines whether the requested content is cached on the second tier caching server (decision block 402). If the requested content is cached in the second tier caching server (Yes in decision block 402), the second tier caching server initiates return of the requested content (act 411) to a client computing system associated with the content request. Referring to FIG. 2, if the second tier caching server 213A had the requested content, the requested content would be returned along the reverse path along arrow 234, arrow 233 and arrow 232 in sequence to the client 201D. While on the return path the requested content may perhaps be cached at the edge caching server 212C. Although the requested content might always be cached in the edge caching servers encountered in the return path, caching policy might also be applied in order to determine whether to cache the returned content.

If the requested content is not cached in the second tier caching server (No in decision block 402), the second tier caching server identifies which of the next tier of caching servers the requested content would be in (act 421) if the requested content were to be cached within the next tier caching server. In this case, the next tier would actually be the origin servers 202. The content request is then distributed to the identified next tier caching server (act 422). For instance, if the requested content were not in the second tier caching server 213A, the content request would be dispatched to the identified next tier caching server. In this case, the origin server 202A may be identified based on the URL itself. The origin server 202A would then return the requested content along the reverse path along arrow 235 (thereupon being potentially cached in second tier caching server 213A), along arrow 234 (thereupon being potentially cached in edge caching server 212C), along arrow 233 and along arrow 232.

Accordingly, an effective mechanism for managing a distributed cache hierarchy is described. The mechanism provides predictability in where the content request should be dispatched, while also providing increased hit rate for caching servers. FIG. 2 also shows an intelligent Domain Name Server (DNS) server 240 to represent that the presence of the principles described with respect to FIGS. 3 and 4 are not inconsistent with the use of an intelligent DNS server. Some of the content requests may still be dispatched to the intelligent DNS server 440 for handling in accordance with the prior art.

The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of

the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed is:

1. A distribution network for retrieving cached content comprising:

a load balancer that receives requests from one or more client computing systems and then routes said content requests to one or more edge routing servers, the load balancer being unaware of where any content corresponding to the content requests may reside in the distribution network;

a plurality of edge caching servers each caching content that is mutually exclusive from the cached content of other caching servers;

one or more edge routing servers hierarchically situated between the load balancer and the plurality of edge caching servers, and wherein the one or more edge routing servers are configured so that they do not independently cache content, said one or more edge routing servers routing content requests received from said load balancer to one of the plurality of edge caching servers; wherein said routing is performed according to a predictive routing protocol which predicts which of the plurality of edge caching servers contains cached content corresponding to given content request by:

an act of identifying which of the plurality of edge caching servers is correlated to the given content request based on the predictive routing protocol, and

an act of routing the content request to the identified edge caching server without having to first determine whether the edge routing server contains the requested content.

2. The network of claim 1, wherein the load balancer routes the content requests based at least in part on a health status associated with each of the one or more edge routing servers, the load balancer being unaware of where any content corresponding to the given content request, resides.

3. The network of claim 2, wherein the load balancer routes the content requests based at least in part on a load balancing algorithm.

4. In a distribution network for retrieving cached content, a hierarchy of servers arranged in tiers for retrieving cached content, the hierarchy of tiers comprising:

a load balancer that receives requests from one or more client computing systems and then routes said content requests to one or more edge routing servers, the load balancer being unaware of where any content corresponding to the content requests may reside in the distribution network;

a plurality of edge caching servers each caching content obtained from one or more origin servers, and the cached content of each edge caching server being mutually exclusive from the cached content of other edge caching servers;

and wherein at least one of the edge caching servers comprises:

an edge caching unit having at least one hardware processor and memory for caching content that is mutually exclusive from the cached content of any other caching unit;

an edge routing unit which receives content requests from the load balancer, and the edge routing unit having at least one hardware processor and memory containing computer-executable instructions which, when implemented, cause the edge routing unit to

9

process one or more content requests received from a load balancer that is unaware of where any content corresponding to the content requests may reside in the distribution network, and the edge routing unit processing the one or more requests according to a predictive routing protocol which predicts where cached content corresponding to requested content resides in the distribution network by:

an act of identifying which of the plurality of edge caching servers is correlated to the requested content based on the predictive routing protocol, and

an act of routing the content request to the identified edge caching server without having to first determine whether the edge routing server contains the requested content.

5. The distribution network in accordance with claim 4, wherein the edge routing unit sends the given content request to the edge caching unit and the edge caching unit determines whether the edge caching unit has the requested content cached.

6. The distribution network in accordance with claim 5, wherein the edge caching unit initiates return of the requested content to a client computing system from which the given content request was sent.

10

7. The distribution network in accordance with claim 4, wherein the edge routing unit sends the given content request to the edge caching unit and the edge caching unit determines that the edge caching unit does not have the requested content, and wherein the method further comprises the following:

an act of the edge caching unit identifying which of the plurality of edge caching servers the given content request would be in; and

an act of distributing the given content request to the identified edge caching server.

8. The distribution network in accordance with claim 4, wherein the edge routing unit sends the given content request to the edge caching unit and the edge caching unit determines that the edge caching unit does not have the requested content, and wherein the method further comprises the following:

an act of the edge caching unit probabilistically identifying a further tier of edge caching servers at which the given content request would likely be cached; and

an act of routing the given content request to one of the edge servers in the identified next tier of edge caching servers.

\* \* \* \* \*